

Studying accreting black holes and neutron stars with time series: beyond the power spectrum

Simon Vaughan^a and Philip Uttley^b

^aDepartment of Physics and Astronomy, University of Leicester, Leicester, LE1 7RH, UK

^bSchool of Physics and Astronomy, University of Southampton, Southampton SO17 1BJ, UK

ABSTRACT

The fluctuating brightness of cosmic X-ray sources, particularly accreting black holes and neutron star systems, has enabled enormous progress in understanding the physics of turbulent accretion flows, the behaviour of matter on the surfaces of neutron stars and improving the evidence for black holes. Most of this progress has been made by analysing and modelling time series data in terms of their power and cross spectra, as will be discussed in other articles in this volume. Recently, attempts have been made to make use of other aspects of the data, by testing for non-linearity, non-Gaussianity, time asymmetry and by examination of higher order Fourier spectra. These projects, which have been made possible by the vast increase in data quality and quantity over the past decade, are the subject of this article.

Keywords: Black Hole, Neutron Star, X-ray Binary, Accretion, Power Spectrum, Time Series, Non-linearity

1. INTRODUCTION

One of the most fundamental problems in modern astrophysics is to understand the flow of matter accreting onto compact objects such as neutron stars and black holes. These systems are powerful and highly variable sources of X-rays, and it is by studying the fluctuations in X-ray luminosity that we have made progress towards understanding the physics involved. These studies are reviewed by Belloni¹ and van der Klis² in this volume. However, almost all this work is dedicated to measuring, modelling and understanding the power spectrum (see Fig. 1) and related second-order statistics (cross spectrum, autocorrelation function, etc.) which necessarily contain information on the statistical moments up to second order only. If the variability processes associated with accretion onto compact objects are linear, Gaussian and stationary then no information is lost (everything is contained in the second order statistics). Until recently this was an implicit and virtually untested assumption, rather than an open area of investigation, which may in part be due to the fact that constructing physical models of the power spectra has proved to be a serious challenge. However, these assumptions cannot be strictly true since we observe high amplitude fluctuations (rms ~ 20 per cent is quite typical) while our knowledge of the physics tells us the X-ray luminosity must be non-negative; the fluctuations in X-ray luminosity cannot be exactly Gaussian if the luminosity is constrained to be always non-negative. Other statistics, that are sensitive to the non-Gaussian and non-linear properties of data, may provide more discriminating tests of existing models or point the way towards new ideas. This article discusses one property of time series from accreting black holes and neutron stars that is not revealed by any individual power spectrum: the rms-flux relation.

2. THE RMS-FLUX RELATION

Put simply, the rms-flux relation describes the fact that the root mean square (rms) amplitude of X-ray variations over some time interval scales linearly with the average X-ray flux (or brightness) during the same time interval. This effect was detected by measuring power spectra for many short segments of a time series (subtracting off any additional power contributed by the measurement errors), using these to calculate the variance over some frequency range $[f_2, f_1]$ by integration (the square root of which is the rms), and comparing these rms values to the mean flux level for each segment*. Due to the stochastic nature of the data (the power spectra are usually

E-mail: sav2@star.le.ac.uk

*The flux and rms measurements have the same units. In X-ray astronomy this is conventionally the detector count rate, i.e. counts s^{-1} .




Figure 1. Power spectrum of the X-ray count rate from an *RXTE* observation of the accreting black hole Cygnus X-1 taken on 16-19 Dec 1996, sampled in bins of $\Delta t = 2^{-8}$ s (3.90625 ms). On the left is shown the power density, revealing the approximately broken power law spectral shape, with a $1/f^0$ spectrum below 0.06 Hz, a roughly $1/f$ spectrum in the range $\sim 0.06 - 6$ Hz, and a $1/f^2$ spectrum at higher frequencies. On the right are exactly the same data but plotted as frequency times power, which better emphasises which frequencies dominate the variance (integrated spectrum). In this $f \times P(f)$ representation a flat spectrum means equal power per logarithmic frequency interval (i.e. $1/f$ spectrum). This plot also better illustrates the slight deviations from power law form. As is conventional the power density is given in fractional units ($[\text{rms}/\text{mean}]^2 \text{ Hz}^{-1}$).

broad-band noise over several decades in frequency; see Ref. 1 and references therein) there is substantial scatter in the individual rms measurements; averaging the rms measurements in flux bins reduces this scatter. The result is the average dependence of rms on ‘local’ flux. When this procedure was applied to data from accreting black holes and neutron stars taken by the *Rossi X-ray Timing Explorer (RXTE)* a strong, linear correlation between rms and flux was revealed³ (see Fig. 2).

The linear rms-flux relation has been observed in time series from stellar-mass black holes and neutron stars,^{3,4} and subsequently confirmed in active galactic nuclei.^{3,5–10} The linear relation appears to be quite ubiquitous in these systems, being present over a range of source ‘states’,⁴ and also occurring in the highly coherent pulsations from an accreting millisecond pulsar.¹¹ The non-zero offset of the relation (Fig. 2) could indicate at least two emission components contribute to the observed emission, one component with the linear dependence between rms and flux, and another with a constant flux, rms or both.

At first sight this relation may seem trivial, merely a particularly simple (“pure”) form of heteroscedasticity. Indeed, the linearity of the rms-flux relation is equivalent to a form of scale invariance. If there was no relation between flux and rms (or a non-linear relation) one could calculate the flux of a given a time series even if no flux scale were given. In the absence of a scale parameter one can still calculate the rms in relative (fractional) units, equivalent to rms/flux, which would be a monotonic function of flux. A given fractional rms would therefore correspond to a unique flux level. Only in the presence of a linear rms-flux relation, however, is the fractional rms independent of flux[†], and so we cannot assign an absolute scale to the time series (flux or rms).

Established tests for heteroscedasticity could have revealed the correlation and shown it to be linear,^{12–14} but the effect was not noticed because of another convention in X-ray astronomy. It is standard practice to plot power spectra in fractional units, i.e. normalise the power density by the squared flux, which removes the dependence on the sensitivity of the detector/telescope used (i.e. in relative units rather than a function of count s^{-1}), but this coincidentally removes the dependence of the power on the flux.

[†]Here we have neglected any small contribution to the total X-ray emission from a constant flux or rms component.




Figure 2. The rms-flux relation. The left panels show the time series data from the Dec 1996 observation of Cyg X-1 with *RXTE* which lasted for two days. The gaps in the coverage are due to observational constraints (the satellite is in a low Earth orbit meaning that every ~ 96 minute orbit the Earth lies between the target source and the satellite). The top panel shows the X-ray flux averaged over 256 s intervals (from data originally binned at 2^{-8} sec resolution). The bottom panel shows the 2 – 20 Hz rms calculated from the integral of the periodogram of each 256 s interval. The two time series clearly show the same trends, indicating the 2 – 20 Hz is correlated with variations in the average flux that occur at much lower frequencies. However, the two are not identical, of note is the increased scatter in the rms values, which is inherent in values estimated from any finite realisation of a stochastic process. The right panel shows a plot of the 2 – 20 Hz rms averaged in flux bins. In this case both the rms and flux were calculated from non-overlapping 1 s intervals, and the rms values were averaged in bins of increasing flux. The binning reduces the effect of the scatter in the rms values, and clearly demonstrates the tight linear dependence of the average rms on the flux.

3. FREQUENCY DEPENDENCE

Before moving on one crucial point needs to be made: the rms-flux relation observed from accreting neutron stars and black holes remains linear on *all* timescales it has been tested on, despite the sometimes complex frequency dependence of the power spectrum. In particular, there are four frequencies (or timescales) that are important when measuring an rms-flux correlation.

Consider a time series x_k of length N sampled at intervals (or averaged over bins) of length Δt , such that the total duration of the data is $N\Delta t$. Now, we can break the time series into consecutive segments of length n , giving $M = N/n$ segments, and for each segment we may calculate the mean level of the time series \hat{x}_i ($i = 1, 2, \dots, M$) and a periodogram $P(f_j) = 2\Delta t |X_i(f_j)|^2/n$ (where X_i is the Fourier transform of the data in segment i , and the factor $2\Delta t/n$ normalises the periodogram such that integrating over positive frequencies gives the variance). By integrating each periodogram over frequencies $[f_2, f_1]$ we get a sample variance for each segment, $\hat{\sigma}_i^2$, over that frequency range. For comparing rms and flux there are four frequencies (or timescales) involved:

- The duration of the time series $f_4 = 1/N\Delta t$
- The length of each segment analysed $f_3 = 1/n\Delta t$
- The lower and upper bounds of the integration f_2, f_1

These are constrained such that $f_4 < f_3 \leq f_2 < f_1$. The constraint that $f_3 \leq f_2$ comes about because we cannot integrate a periodogram over frequencies lower than are measured (i.e. $f_2 \geq 1/n\Delta t$), and similarly $f_1 \leq 1/2\Delta t$, which means we cannot integrate over frequencies above the Nyquist limit. Frequencies f_4 and f_3 determine range of frequencies covered by the time series of mean values, \hat{x}_i , and rms values, $\hat{\sigma}_i$, while the frequencies f_2, f_1 determine over what range of frequencies the rms is accounting.

The observed rms-flux relation means that changes in X-ray flux occurring over the frequencies range $[f_4, f_3]$ are correlated with the simultaneous rms measured from the frequency range $[f_2, f_1]$. The relation appears linear irrespective of the values of f_4, f_3 and f_2, f_1 used. It is this fact that makes the relation interesting. We could, for example, take a time series from a stationary, linear, Gaussian process which has a constant rms, and multiply it by a steep linear function. The effect would be to increase both the mean level and the rms where the linear function is highest, and decrease both where the linear function is lowest, so as to produce a linear relationship between rms and (mean) flux. However, on timescales over which the linear function varies by only a small amount, where the variance is dominated by variations in the original series, there will be no correlation between rms and flux, by construction. In this case the relation does not exist on all timescales.

The power spectrum for the accreting black hole system Cygnus X-1 is shown in Fig. 1, revealing broad-band noise that has a red ($1/f$ or steeper) spectrum above ~ 0.06 Hz (~ 16 -s) and is roughly white below this (although the power spectrum may steepen again at substantially lower frequencies). Therefore, time series segments substantially shorter than 10 s are red noise, with the variance (power) dominated by longest timescales (lowest frequencies), and are weakly non-stationary (the mean can only be well defined on much longer timescales), even though they are all realisations of the same stochastic process. Nevertheless, regardless of the timescales used to measure the mean ‘local’ flux, and the frequencies used to measure the rms, the two are linearly related (see Fig. 1). As shown by Gleissner et al.,⁴ the relation remains present in data from Cygnus X-1 even when measured on timescales of months-years. This point is crucial for understanding the origin of the relation.

4. A SIMPLE PHENOMENOLOGICAL MODEL

A linear, Gaussian, stationary, stochastic processes can be constructed by addition of independent signals in either the time domain (e.g. a moving average [MA] or shot noise process) or the Fourier domain (by the sum of Fourier components with independent phases). In this scenario there need be no coupling between variations at different frequencies – high frequency variations require no knowledge of the behaviour of lower frequency variations. This is not true for a red noise process with a linear rms-flux relation on all timescales, which does require higher frequency variations to respond to the activity of all lower frequency variations. The power spectrum is red noise over most of the available bandpass and so the mean flux in any short time series segment is determined by the long timescale variations (which dominate the overall variance), yet the rms on much shorter timescales is correlated with this. Further demonstrations of this are given in Refs. 3, 15. This suggests a ‘top-down’ process is required, in which longer timescale (lower frequency) variations originate first and modulate shorter timescale (higher frequency) variations. In fact what is required is simple amplitude modulation whereby the amplitude of variations at any given frequency is modulated by the variations at all lower frequencies.

This suggests a simple way to synthesise a stochastic time series with the linear rms-flux relation on all timescales. A general way to synthesise any linear, Gaussian, stochastic time series is by the summation, in the time domain, of individual sinusoidal signals with (independent) random phases.

$$y_k = \sum_{j=1}^{N/2} A_j \sin(2\pi f_j t_k + \phi_j) = \sum_{j=1}^{N/2} a_j(t_k) \quad (1)$$

where $y_k = y(t_k)$ is the output time series and A_j and ϕ_j are the amplitude and phase (uniform $\phi_j \in [0, 2\pi)$) at frequency f_j . This is simplified by writing $a_j(t_k) = A_j \sin(2\pi f_j t_k + \phi_j)$. For a realistic signal with finite variance $A_j \rightarrow 0$ as $N \rightarrow \infty$, and hence $a_j \rightarrow 0$. In order to impose a linear rms-flux relation on the output time series we simply replace the sum with a product. By multiplying the components (again, in the time domain) we ensure that variations at any given frequency are modulated by variations on all other frequencies:

$$x_k = \prod_{j=1}^{N/2} \{1 + A_j \sin(2\pi f_j t_k + \phi_j)\} \quad (2)$$

with $0 \leq A_j < 1$ to ensure the output signal is positive. Consider a simple signal comprising two sinusoids, one low frequency, one high frequency. The product of these gives an output signal where the (rms) amplitude on




Figure 3. Exponentially transformed simulated data. The left panels shows two versions of an artificially generated time series in which the data are the same but the amplitudes (absolute rms levels) were varied prior to exponentially transforming. After applying the exponential transform the two series have different fractional rms values. The data were generated using the power spectrum shown in Fig. 1. The right panel illustrates the spectral distortion caused by the transformation. The solid curve shows the Lorentzian power spectrum of an input process (linear) and the data show the estimated power spectrum for artificial data after exponential transformation. Note the logarithmic scale that clearly reveals the weak distortions introduced by the transformation, including a harmonic.

short timescales (dominated by the high frequency sinusoid) scales with the level of the low frequency sinusoid, which dominates the mean level, hence produces a linear rms-flux relation. Extending this procedure we can generate a signal with a linear rms-flux relation on all timescales by multiplying sinusoids from all frequencies (for a power spectrum defined by the A_i values).

It is interesting to compare the properties of the additive and multiplicative methods. Taking the logarithm of equation 2 we get:

$$\log[x_k] = \sum_{j=1}^{N/2} \log[1 + a_j(t_k)] \approx \sum_{i=1}^{N/2} a_j(t_k) = y_k \quad (3)$$

where we have used the Taylor series expansion of $\log[1 + a] = a - a^2/2 + a^3/3 + \dots$ (noting that $|a_j(t_k)| \ll 1$), and kept only the first order term (the higher order terms in the Taylor series may be safely neglected, see Ref. 15 for more detail). This result indicates that taking the logarithm of the signal generated by the multiplicative method (equation 2) gives the same result as the additive method (equation 1). Reversing this argument we see that to impose a linear rms-flux relation on an otherwise Gaussian signal we must only make an exponential transformation: $x_k \propto \exp[y_k]$. (This fact could have been obtained from the known variance-stabilising methods,^{12–14} except that it is not immediately obvious how these methods would be affected by the non-white power spectrum, i.e. auto-correlation, of the data under consideration here.)

5. SIMULATIONS OF EXPONENTIALLY TRANSFORMED DATA

Now that we have a simple model for generating time series by a simple transformation of linear data, we are in a position to produce artificial data that obey the rms-flux relation simply by generating linear, Gaussian data using standard techniques, and the applying the static exponential transformation. In this way we can very easily generate data with a range of properties (mean, variance, power spectral shape) and a linear rms-flux relation, then observe the consequences of the relation.

One of the most widely-used and computationally efficient methods for generating linear, Gaussian time series with arbitrary power spectrum $P(f)$ is the Davies-Harte method,¹⁶ where $\sqrt{P(f)}$ is used to define the data in the Fourier domain, and the real and complex components of the Fourier transform are each randomised

by multiplying by independent normal deviates[‡]. The data are then inverse Fourier transformed to the time domain using an FFT. This is a computationally efficient method for realising the summation of equation 1 (with random, uniform phases $\phi_j \in [0, 2\pi)$ and amplitudes distributed as $A_j^2 \sim P(f_j)\chi_2^2/2$, where χ_2^2 is a random deviate distributed as chi-square with two degrees of freedom).

Fig. 3 shows short sections of simulated times series produced by exponentially transforming data produced with the Davies-Harte method and a power spectrum taken from Fig. 1. Exactly the same random number sequence was used to generate the three series, the only difference being the total variance (the normalisation of the power spectrum model). These clearly demonstrate the positively skewed flux distribution and how this becomes more skewed as the total variance increases.

The exponential transformation modifies the mean and variance of the data, therefore we must know the conversion between linear and transformed data in order to generate transformed data with exactly specified statistical properties. The relations between linear and exponential data for the mean and variance is the same as those relating the Gaussian and log-normal distributions. Specifically, for a linear data y_k with zero-mean ($\mu_y = 0$) and variance σ_y^2 , the mean and variance of the exponentially transformed x_k data are:

$$\mu_x = \exp[\sigma_y^2/2] \quad \sigma_x = \exp[\sigma_y](\exp[\sigma_y] - 1) \quad (4)$$

The power spectrum of the transformed data is also not identical to that of the input linear data (see Fig. 3). For broad-band spectra the distortion is small, but becomes strongest for strongly peaked spectra at high amplitudes. The effect of this distortion is described by Johnson¹⁸ (section III-D), who gives a formula for the effect of transformation on the auto-covariance function. This can be inverted to give the formula for the input auto-covariance that produces the desired output auto-covariance in the transformed data.

The above recipe can produce synthetic time series with the linear rms-flux relation with a wide range of power spectra. However, some processes, particularly those with very high power concentrated in narrow-band features have auto-covariances whose Fourier transforms contain negative values. This means they have negative powers, which cannot correspond to a real-valued time series, meaning no time series can be generated with the Davies-Harte method. A time series with a strong, narrow-band spectrum (e.g. a narrow Lorentzian) can be exponentially transformed but the resulting time series has a spectrum with harmonics. If we try to suppress these harmonics by using the Johnson formula to specify the spectrum of the transformed data as a single, narrow Lorentzian, we find that the recipe fails – un-distorting the auto-covariance produces a function that has negative-valued Fourier transform, i.e. negative powers, that cannot be realised in a time series. This may indicate that it is not possible for Nature to produce a strong, narrow band signal without additional harmonics if the linear rms-flux relation is present (Fig. 3).

6. CONSEQUENCES OF THE EXPONENTIAL TRANSFORMATION

If we consider the linear rms-flux relation in purely signal processing terms, it can be seen as an effect of the exponential transformation of a linear, Gaussian signal (e.g. the output of equation 1). This has several consequences.

Log-normality: The linear rms-flux relation is equivalent to an exponential transformation of Gaussian data, which in turn predicts a log-normal distribution of fluxes (over long timescales), since the exponential transformation of a Gaussian distribution is log-normal. The long observation of Cygnus X-1 taken in December 1996 provides a stunning confirmation of this prediction (Fig. 4; Ref. 15). However, one must again be careful to account for the weak non-stationarity of signals with red power spectra, for which the flux distribution is dominated by the lowest frequency variations, which are not sufficiently sampled to converge on the correct underlying distribution. Only by measuring the distribution over sufficiently long timescales that the power spectrum is nearly flat (white) can we be sure to completely sample the true distribution.

[‡]Strictly speaking, the method described by Davies & Harte¹⁶ begins by defining the process in term of its auto-covariance, then taking the Fourier transform to obtain the power spectrum $P(f)$. The method of generating linear, Gaussian time series by randomising the Fourier transform and transforming back to the time domain using an FFT was introduced to the astronomical literature by Timmer & König.¹⁷




Figure 4. The left panel shows an X-ray time series from Cygnus X-1. Specifically, this is a 20 s interval from the start of the much longer Dec 1996 observation, and for clarity the data were rebinned on a scale of 0.0625 s (2^{-4} s) from data originally sampled at 2^{-8} s. Compare with the simulations of Fig. 3. The right panel shows the empirical density function for X-ray fluxes for Cygnus X-1 from December 1996. The histogram (with error bars) shows the empirical density and the solid grey curve is the best-fitting lognormal distribution, and the bottom panel shows the normalised data-model residuals from this fit. The lognormal fit to the histogram gave $\chi^2 = 238.6$ for 189 degrees of freedom, not a great fit, but good considering the very high signal-to-noise. The empirical density is the normalised histogram of fluxes averaged over 0.25 s long time intervals. For each continuous segment of data (~ 3000 s per *RXTE* orbit) the fluxes at 0.25 s intervals were normalised to the mean level for that particular data segment – this reduces any slight shift in the mean of the distribution due to very low frequency power (or non-stationary trends). This gave a total of 358,524 flux samples that were used to produce the histogram.

Non-linearity: Another consequence of the exponential transformation model is that time series showing a linear rms-flux relation on all timescales are, in one sense, non-linear[§]. Any linear, stochastic model may be written as an MA. In discrete notation:

$$y_k = \sum_{i=0}^{\infty} g_i \varepsilon_{k-i} \quad (5)$$

where ε_k are independent and identically distributed (*iid*) random deviates, and g_i is the kernel that defines the ‘memory’ of the process (see Ref. 19 for details). A non-linear model cannot be written in this form but requires higher order terms. With some generality, a non-linear model may be written as a Volterra series:¹⁹

$$x_k = \sum_{a=0}^{\infty} G_a \varepsilon_{k-a} + \sum_{a=0}^{\infty} \sum_{b=0}^{\infty} G_{ab} \varepsilon_{k-a} \varepsilon_{k-b} + \sum_{a=0}^{\infty} \sum_{b=0}^{\infty} \sum_{c=0}^{\infty} G_{abc} \varepsilon_{k-a} \varepsilon_{k-b} \varepsilon_{k-c} + \dots \quad (6)$$

The model is linear if the terms G_{ab}, G_{abc}, \dots are all zero (in which case it reduces to equation 5). This is not true if we consider the exponential transformation of a general linear process:

$$x_k = \exp[y_k] = 1 + y_k + \frac{(y_k)^2}{2} + \sum_{n=3}^{\infty} \frac{(y_k)^n}{n!} = 1 + \sum_{i=0}^{\infty} g_i \varepsilon_{k-i} + \frac{1}{2} \left(\sum_{i=0}^{\infty} g_i \varepsilon_{k-i} \right)^2 + \sum_{n=3}^{\infty} \frac{1}{n!} \left(\sum_{i=0}^{\infty} g_i \varepsilon_{k-i} \right)^n \quad (7)$$

The last equality made use of the expression of a general linear model (equation 5). This may be re-arranged to form a Volterra series:

$$x_k = 1 + \sum_{a=0}^{\infty} G_a \varepsilon_{k-a} + \sum_{a=0}^{\infty} \sum_{b=0}^{\infty} G_{ab} \varepsilon_{k-a} \varepsilon_{k-b} + \sum_{a=0}^{\infty} \sum_{b=0}^{\infty} \sum_{c=0}^{\infty} G_{abc} \varepsilon_{k-a} \varepsilon_{k-b} \varepsilon_{k-c} + \dots \quad (8)$$

[§]Here we should note that only models for processes can be non-linear, and not time series data themselves, strictly speaking. What we really mean is that this simple phenomenological model that reproduces the main features of these time series (power spectrum shape and linear rms-flux relation) is non-linear.

with the coefficients $G_a \propto g_a$, $G_{ab} \propto g_a g_b$, $G_{abc} \propto g_a g_b g_c$, and so on. These higher order terms are non-zero if g_i is non-zero for more than one i (i.e. the power spectrum is not white), meaning the exponential transform of a general linear process cannot (in general) be written in terms of a linear model. This argument is discussed in more detail by Ref. 15.

The fact that accreting compact objects do show such a linear rms-flux relation on all observable timescales may mean that their time series can be modelled in terms of a linear, Gaussian process subject only to an exponential transformation, which after all is a static non-linear transformation, and suggests there is no need for dynamical non-linearities.

Millisecond flares: Cygnus X-1, and perhaps other X-ray binaries, occasionally display strong flares, bursts of increased X-ray activity lasting ~ 0.1 seconds. These have been identified as being inconsistent with the variability arising from a stationary, Gaussian process,²⁰ but have been shown to be consistent with the output of an exponentially transformed Gaussian process.¹⁵

Prolonged ‘low flux’ episodes in AGN: Some of the most well-studied Active Galactic Nuclei (AGN; see Ref. 21 in this volume) display episodes of very low activity lasting for days-months, when the X-ray flux remains low and relatively stable.²² These can be easily understood as a consequence of the linear rms-flux relation. At high fluxes the variance is also high, and so the source appears very X-ray ‘active’, while at much lower fluxes the variance is also at a minimum, making the source appear to be in a relatively stable ‘low’ state (see Fig. 3). The difference between ‘high’ and ‘low’ activity periods becomes more exaggerated as the gradient of the rms-flux relation increases. This gradient is equivalent to the fractional rms - those sources with strongest overall variability are also those with the most marked transitions from ‘active’ to ‘inactive’ periods (on an absolute scale).

7. PHYSICALLY MOTIVATED MODELS

The presence of a linear rms-flux relation on all timescales has some important consequences for our understanding of the physics involved in producing variable X-ray emission from accreting compact objects. The relation appears the same in all known accreting black hole systems (stellar mass¹ and supermassive²¹ black holes), and also in at least one accreting neutron star system.¹¹ As such it seems rather general to accreting compact objects, and any physical model of the variability mechanism must reproduce not only the correct power spectrum but also the rms-flux relation in order to be acceptable. One important point that can be made immediately is that models involving dynamical non-linearities (chaos) are not required: The non-linear behaviour of accreting systems can be explained by a simple, static non-linear transformation, eliminating the need for dynamical non-linearities.

7.1. Models that fail the rms-flux test

Almost since the discovery of randomly fluctuating X-ray brightness from Cygnus X-1 the most popular models for the X-ray variability have all been variants of shot noise.²³ By construction, shot noise can reproduce any power spectral shape, in the same way that an MA process can have any conceivable power spectrum. Therefore a shot noise model is always capable of matching the power spectra of time series from accreting compact objects, and it is quite straightforward to imagine individual and independently occurring shots as the result of ‘active regions’ appearing and disappearing around or on an accretion disc in a manner not dissimilar to X-ray emitting flares above the surface of the Sun.

But while shot noise can reproduce any power spectrum, it also predicts a roughly Poisson or Gaussian distribution of X-ray fluxes. This is because the observed time series do not appear as a series of clearly resolved events or shots but as a fluctuating continuum, as is reflected in the very broad power spectrum. In the limit of a large number of independent shots occurring at any one time the output time series will be linear and Gaussian (by the central limit theorem) and will not show any dependence of rms on flux nor a lognormal flux distribution.

A closely related model is Self-Organised Criticality^{24–26} (SOC), where certain systems of interacting elements reach a critical point at which events occur on all possible scales with a power law distribution. Indeed, there have been attempts to model X-ray emission from accretion discs with toy models based on SOC.^{27, 28} But, as pointed out by Bak et al.²⁵ (see also section 4 of Christensen et al.²⁶), SOC produces power law distributions of event scales. These events need to be randomly superposed to form shot noise with a broad power spectrum




Figure 5. The right panel shows power spectra for two time series derived from the same Cygnus X-1 data: (black) the time series of X-ray fluxes (as Figure. 1) and (grey) a time series of 4 – 32 Hz rms amplitudes. Both were calculated using the same 0.25 second segments. Clearly the spectrum of the rms is very similar to the spectrum of the flux, except for at the highest frequencies. (The estimated spectra from both the flux and rms data contain flat ‘background’ noise terms that must be subtracted to reveal the spectrum of the true process. In the former case the background noise is due to Poisson noise in the collection of X-ray counts by the detector, in the latter case it is due to random variations in the variance inherent in any finite realisation of a stochastic process.) The right panel shows the results of the same analysis applied to artificial data. The solid curve shows the power spectrum used to produce linear, Gaussian data (from the measured power spectrum of the real Cyg X-1 data). Time series generated from this spectrum were exponentially transformed, and the mean flux and rms measured from segments of length 0.25 s, exactly as for the real data. The effect of the transformation, and also the averaging, slightly distorts the spectrum of the mean fluxes away from the original input spectrum. But it is clear the power spectrum of the rms is similar between the real data and the transformed artificial data, indicating that the exponential transformation is largely (but perhaps not entirely) responsible for the structure in the spectrum of the rms.

(e.g. $\sim 1/f$). By the central limit theorem, the resulting flux distribution will be Gaussian again. SOC is a model for generating a power law distribution of event scales, not a model for interaction of events.

7.2. Models that pass the rms-flux test

As described above, a model involving the modulation (multiplication) of independent fluctuations, rather than the superposition (addition), can reproduce the linear rms-flux relation. One model of accretion flows that has this property was discussed by Lyubarskii.²⁹ In this model random fluctuations in the accretion rate occur at all radii and propagate through the disc. The X-ray emission is produced only from the inner edge of the disc, near the compact object (black hole or neutron star), and responds in proportion to the total mass accretion rate at the inner edge, which is the product of the fluctuations propagating in from all radii. Longer timescale (lower frequency) fluctuations are generated in the outer regions and propagate inwards modulating the more rapid fluctuations generated in inner regions. This model has been developed further^{30,31} and can reproduce many of the observed timing properties including the linear rms-flux relation. Another model that also produces amplitude modulated fluctuations was discussed by King et al.³²

8. OTHER DIAGNOSTICS

If the rms of a signal, calculated over any given frequency range, is linearly correlated with the signal’s strength, then the rms values should form a time series with the same power spectrum as that of the signal itself. This can be tested by estimating the power spectrum of the rms time series. There is a complication due to the scatter in rms values caused by the stochastic nature of the underlying process. However, if the power spectrum of the process is known (or can be estimated accurately, as is the case here), then the amplitude of these random fluctuations on the rms can be calculated, and its effect of the power spectrum of the rms can be

subtracted. Fig. 5 shows the power spectrum of the signal (source flux) compared to the power spectrum of the rms (computed over 4 – 32 Hz). Clearly the two are very close, but not identical. The slight difference in normalisation at lower frequencies is mostly due to the presence of Poisson noise from the detection process, which adds an approximately constant amount to the rms (and therefore lowers the fractional amplitude of the frequency dependent variations). The divergence at higher frequencies (> 1 Hz) is more intriguing. The right panel shows the results of the same procedure applied to data simulated using the observed power spectrum, and transformed using the exponential transform. As predicted the power spectrum for the signal and rms match closely. The difference between the real data and simulated data may indicate the coupling between flux and rms is not quite as simple as predicted by the exponential transformation. (Another technique not yet applied to X-ray data is to calculate the coherence between two time series, one for the mean flux and one for the rms, to better understand their coupling.)

Another potentially powerful diagnostic of frequency coupling is the bispectrum, or bicoherence. Where the power spectrum is a second-order spectrum, formed from the product of two Fourier transforms $[X(f_j) X^*(f_j)]$, the bispectrum is a third-order spectrum formed from the product of three Fourier transforms $[X(f_j) X(f_k) X^*(f_j + f_k)]$. The computation and interpretation of bispectra (and higher-order spectra in general) is treated by Refs 33,34, and they have been applied to data from various fields including geology,³⁵ plasma physics³⁶ and speech processing.³⁷

The bispectrum is a complex quantity defined for the pair of frequencies (f_j, f_k) . The phase of the bispectrum is often called the *biphase* and its modulus-squared amplitude is called the *bicoherence* (once it has been suitably normalised). The bispectrum provides a measure of the multiplicative nonlinear interaction of frequency components in the time series.³⁸ A simple way to understand the bicoherence is as follows. If the variances occurring at the frequencies $f_j, f_k, f_j + f_k$ are independent (spontaneously excited fluctuations) they will have independent, random phases. The bispectrum will thus tend to zero. If there are variations occurring at $f_j + f_k$ that are due to multiplicative coupling of variations at f_j and f_k then there will be phase coupling and the bispectrum will be non-zero. The value of the bicoherence in some sense represents the fraction of variability power due to multiplicative coupling. A stationary, linear, Gaussian processes has a bispectrum that vanishes (at all frequencies in the large sample limit), and therefore the amplitude of the bicoherence is (asymptotically) zero. For stationary, linear, non-Gaussian process bicoherence is a non-zero constant while for a stationary non-linear process the bicoherence is not constant with frequency. These properties make the bispectrum sensitive to the non-Gaussian and non-linear characteristics of the time series.³⁸

The success of the amplitude modulation description of the Cygnus X-1 data suggests that fluctuations at different frequencies are coupled, meaning the rms-flux relation should accompany structure in the bicoherence. Fig. 6 clearly reveals this to be true; the non-trivial structure in the bicoherence indicates non-linear frequency coupling. However, interpreting the bicoherence is not simple. The bicoherences at each bifrequency (f_j, f_k) are not independently distributed, making it difficult to quantify error bars and therefore apply model-fitting procedures. Furthermore, the presence of Poisson noise in these data causes spurious bicoherence at high frequencies because while the noise is uncorrelated with the signal (i.e. the true source flux), it is not independent, since the amplitude of Poisson fluctuations depends on $\sim \sqrt{flux}$. At high frequencies the Poisson noise dominates over the intrinsic variations in the source flux (which drop off rapidly with increasing frequency). Nevertheless, bicoherence may become a powerful test for more advanced models of variability from black holes.^{15, 39, 40}

9. ANALOGUES IN NATURE

Apart from accreting compact objects, do other natural processes that have broad band noise spectra and also show a linear rms-flux relation, or is this something peculiar to the dynamics of accretion flows?

Zhang⁴¹ has made comparisons with Solar flares and the flaring behaviour of Gamma-Ray Bursts (GRBs), but these are transient events and do not have well-defined broad band noise spectra. As a result of this, it is not possible to conclude whether a linear rms-flux relation occurs on all timescales (fluctuations at all frequencies modulate each other) or whether it occurs because there is an overlying exponential decay that modulates small scale flickering. For example, we could produce a time series using linear shot noise and then multiply this by a strong exponentially decaying trend. The trend dominates the average level and also modulates the variance from




Figure 6. Bicoherence for the Cygnus X-1 data. The left panel shows the bicoherence from the ‘principal domain’ of bi-frequencies f_j, f_k (the principal domain is defined as those pairs of frequencies for which f_j spans the full range of Fourier frequencies $[0, 1/2\Delta t]$, with the additional constraints $f_k \leq f_j$ and $f_j + f_k \leq 1/2\Delta t$). The effect of detector noise has been subtracted (to first order). The bicoherences were binned in equal logarithmic frequency intervals along each axis, and colour coded such that lighter shades correspond to higher bicoherences. In order to better illustrate the structure in the bicoherence the right panel shows the bicoherence compressed into one dimension as a function of $f_j + f_k$ (by summing over lines of constant $f_j + f_k$). The error bars come from the scatter of individual bicoherences within each logarithmic frequency bin, but as individual bicoherence estimates are not independent at different bifrequencies, these should be taken as illustrative only.

fluctuations at higher frequencies. But beyond this, there is no further modulation by fluctuations of different frequencies. Exactly because of this many economic time series, such as the historical value of the S&P 500 share index, appear at first sight to display the linear rms-flux relation. However, this is the result of inflation – a dominating, long-term increase in the scale. It is not clear whether random fluctuations on shorter timescales modulate random fluctuations on yet shorter timescales.

Some natural phenomena that are known to possess a strong coupling between signal strength and rms. Argollo de Menezes and Barabási⁴² discuss how various the time-dependent traffic on internet routers of the Mid-Atlantic Crossroads network, and stream-flow on rivers in the U.S. river basin, both obey a relation such that the rms of the activity of a sub-system (router or river) scales with the mean level of activity of that subsystem. The scaling is a power law, but in some cases the index is unity, giving a linear rms-flux relation. This is similar to “Taylor’s law” in population dynamics⁴³ and is a well known property of many ecological⁴⁴ and financial systems. The difference between these results and those for accreting X-ray sources is that in the case of the networks, the relation describes how mean and rms values over different subsystems (be they species, routers or rivers) are related, whereas for the X-ray sources we are looking a mean and rms values at different times for just one system. (Although the same relation occurs in many systems.) Despite this difference it may still be possible to compare the characteristics and interpretation of rms-flux relations in these very different systems.

ACKNOWLEDGMENTS

The authors wish to acknowledge financial support through the UK Science and Technology Facilities Council (STFC).

REFERENCES

1. T. Belloni, “Noise components from black-hole binaries in our galaxy,” in *Noise and Fluctuations in Photonics, Quantum Optics, and Communications*, L. Cohen, ed., *Proc. SPIE* **6603**, 2007 in press.

2. M. van der Klis, "Millisecond phenomena in mass accreting neutron stars," in *Noise and Fluctuations in Photonics, Quantum Optics, and Communications*, L. Cohen, ed., *Proc. SPIE* **6603**, 2007 in press.
3. P. Uttley and I. M. McHardy, "The flux-dependent amplitude of broadband noise variability in X-ray binaries and active galaxies," *MNRAS* **323**, pp. L26–L30, 2001.
4. T. Gleissner, J. Wilms, K. Pottschmidt, P. Uttley, M. A. Nowak, and R. Staubert, "Long term variability of Cyg X-1. II. The rms-flux relation," *A&A* **414**, pp. 1091–1104, 2004.
5. R. Edelson, T. J. Turner, K. Pounds, S. Vaughan, A. Markowitz, H. Marshall, P. Dobbie, and R. Warwick, "X-Ray Spectral Variability and Rapid Variability of the Soft X-Ray Spectrum Seyfert 1 Galaxies Arakelian 564 and Ton S180," *ApJ* **568**, pp. 610–626, Apr. 2002.
6. S. Vaughan, A. C. Fabian, and K. Nandra, "X-ray continuum variability of MCG-6-30-15," *MNRAS* **339**, pp. 1237–1255, 2003.
7. S. Vaughan, R. Edelson, R. S. Warwick, and P. Uttley, "On characterizing the variability properties of X-ray light curves from active galaxies," *MNRAS* **345**, pp. 1271–1284, 2003.
8. C. M. Gaskell, "Lognormal X-Ray Flux Variations in an Extreme Narrow-Line Seyfert 1 Galaxy," *ApJ* **612**, pp. L21–L24, 2004.
9. I. M. McHardy, I. E. Papadakis, P. Uttley, M. J. Page, and K. O. Mason, "Combined long and short time-scale X-ray variability of NGC 4051 with RXTE and XMM-Newton," *MNRAS* **348**, pp. 783–801, Mar. 2004.
10. S. Vaughan, K. Iwasawa, A. C. Fabian, and K. Hayashida, "The exceptional X-ray variability of the dwarf Seyfert nucleus NGC 4395," *MNRAS* **356**, pp. 524–530, 2005.
11. P. Uttley, "SAX J1808.4-3658 and the origin of X-ray variability in X-ray binaries and active galactic nuclei," *MNRAS* **347**, pp. L61–L65, 2004.
12. M. S. Bartlett and D. G. Kendall, "Statistical analysis of variance-heterogeneity and the logarithmic transformation," *J. R. Stat. Soc. Ser. B* **8**, pp. 128–138, 1946.
13. M. S. Bartlett, "The use of transformations," *Biometrics* **3**, pp. 39–53, 1947.
14. G. E. P. Box and D. R. Cox, "An analysis of transformations," *J. R. Stat. Soc. Ser. B* **26**, pp. 211–252, 1964.
15. P. Uttley, I. M. McHardy, and S. Vaughan, "Non-linear X-ray variability in X-ray binaries and active galaxies," *MNRAS* **359**, pp. 345–362, 2005.
16. R. B. Davies and D. S. Harte, "Tests for the hurst effect," *Biometrika* **74**, pp. 95–101, 1987.
17. J. Timmer and M. König, "On generating power law noise," *A&A* **300**, pp. 707–710, 1995.
18. G. E. Johnson, "Construction of particular random processes," *Proc. IEEE* **82**, pp. 270–285, 1994.
19. M. B. Priestley, *Spectral Analysis and Time Series*, Academic Press, London, 1981.
20. M. Gierliński and A. A. Zdziarski, "Discovery of powerful millisecond flares from Cygnus X-1," *MNRAS* **343**, pp. L84–L88, 2003.
21. P. Uttley, "Timing the x-ray fluctuations from supermassive black holes," in *Noise and Fluctuations in Photonics, Quantum Optics, and Communications*, L. Cohen, ed., *Proc. SPIE* **6603**, in press.
22. P. Uttley, I. M. McHardy, I. E. Papadakis, M. Guainazzi, and A. Fruscione, "The swan song in context: long-time-scale X-ray variability of NGC 4051," *MNRAS* **307**, pp. L6–L10, 1999.
23. N. J. J. Terrell, "Shot-Noise Character of Cygnus X-1 Pulsations," *ApJ* **174**, pp. L35–L41, May 1972.
24. P. Bak, C. Tang, and K. Wiesenfeld, "Self-organized criticality - An explanation of 1/f noise," *Physical Review Letters* **59**, pp. 381–384, July 1987.
25. P. Bak, C. Tang, and K. Wiesenfeld, "Self-organized criticality," *Phys. Rev. A* **38**, pp. 364–374, 1988.
26. K. Christensen, Z. Olami, and P. Bak, "Deterministic 1/f noise in nonconservative models of self-organized criticality," *Physical Review Letters* **68**, pp. 2417–2420, 1992.
27. S. Mineshige, N. B. Ouchi, and H. Nishimori, "On the generation of 1/f fluctuations in X-rays from black-hole objects," *PASJ* **46**, pp. 97–105, 1994.
28. M. Takeuchi, S. Mineshige, and H. Negoro, "X-Ray Fluctuations from Black-Hole Objects and Self Organization of Accretion Disks," *PASJ* **47**, pp. 617–627, 1995.
29. Y. E. Lyubarskii, "Flicker noise in accretion discs," *MNRAS* **292**, pp. 679–685, 1997.

30. E. Churazov, M. Gilfanov, and M. Revnivtsev, "Soft state of Cygnus X-1: stable disc and unstable corona," *MNRAS* **321**, pp. 759–766, 2001.
31. O. Kotov, E. Churazov, and M. Gilfanov, "On the X-ray time-lags in the black hole candidates," *MNRAS* **327**, pp. 799–807, 2001.
32. A. R. King, J. E. Pringle, R. G. West, and M. Livio, "Variability in black hole accretion discs," *MNRAS* **348**, pp. 111–122, 2004.
33. D. R. Brillinger and M. Rosenblatt, "Computation and interpretation of the k-th order spectra," in *Spectral Analysis of Time Series*, B. Harris, ed., pp. 189–232, Wiley, New York, 1967.
34. C. L. Nikias and A. P. Petropulu, *Higher-order spectra analysis - A nonlinear signal processing framework*, Prentice-Hall, Englewood Cliffs, NJ, 1993.
35. J. Rial and C. Anaclerio, "Understanding nonlinear responses of the climate system to orbital forcing," *Quaternary Science Reviews* **19**, pp. 1709–1722, 2000.
36. Y. C. Kim and E. J. Powers, "Digital bispectral analysis and its applications to nonlinear wave interactions," *IEEE Transactions on Plasma Science* **7**, pp. 120–131, June 1979.
37. J. Fackrell and S. McLaughlin, "Detecting nonlinearities in speech sounds using the bicoherence," *Proceedings of the Institute of Acoustics* **18(9)**, pp. 123–130, 1996.
38. M. J. Hinich, "Testing for gaussianity and linearity of a stationary time series," *Journal of Time Series Analysis* **3**, pp. 169–176, 1982.
39. T. J. Maccarone and P. S. Coppi, "Higher order variability properties of accreting black holes," *MNRAS* **336**, pp. 817–825, 2002.
40. T. J. Maccarone and J. D. Schnittman, "The bicoherence as a diagnostic for models of high-frequency quasi-periodic oscillations," *MNRAS* **357**, pp. 12–16, Feb. 2005.
41. S. N. Zhang, "Similar phenomena at different scales: black holes, the sun, gamma-ray bursts, supernovae, galaxies and galaxy clusters," in *Highlights of Astronomy*, K. A. van der Hucht, ed., *Proc. IAU* **14**, 2007 in press.
42. M. Argollo de Menezes and A.-L. Barabási, "Separating Internal and External Dynamics of Complex Systems," *Physical Review Letters* **93**, pp. 068701–+, Aug. 2004.
43. L. R. Taylor, "Aggregation, Variance and the Mean," *Nature* **189**, pp. 732–735, 1961.
44. B. H. McArdle and M. J. Anderson, "Variance heterogeneity, transformations, and models of species abundance: a cautionary tale," *Can. J. Fish. Aquat. Sci.* **61**, pp. 1294–1302, 2004.